arXiv:1611.03916v1 [cond-mat.other] 12 Nov 2016

A KIM-compliant *potfit* for fitting sloppy interatomic potentials: Application to the EDIP model for silicon

Mingjian Wen[1], Junhao Li[2], Peter Brommer[3], Ryan S. Elliott[1], James P. Sethna[2] and Ellad B. Tadmor[1] ‡

E-mail: tadmor@umn.edu
[1] Department of Aerospace Engineering and Mechanics, University of Minnesota, Minneapolis, MN 55455, USA
[2] Laboratory of Atomic and Solid State Physics, Cornell University, Ithaca, New York 14853, USA
[3] Warwick Centre for Predictive Modelling, School of Engineering, and Centre for Scientific Computing, University of Warwick, Coventry CV4 7AL, UK

Abstract. Fitted interatomic potentials are widely used in atomistic simulations thanks to their ability to compute the energy and forces on atoms quickly. However, the simulation results crucially depend on the quality of the potential being used. Force matching is a method aimed at constructing reliable and transferable interatomic potentials by matching the forces computed by the potential as closely as possible, with those obtained from first principles calculations. The *potfit* program is an implementation of the force-matching method that optimizes the potential parameters using a global minimization algorithm followed by a local minimization polish. We extended *potfit* in two ways. First, we adapted the code to be compliant with the KIM Application Programming Interface (API) standard (part of the Knowledgebase of Interatomic Models Project). This makes it possible to use *potfit* to fit many KIM potential models, not just those prebuilt into the *potfit* code. Second, we incorporated the geodesic Levenberg–Marquardt (LM) minimization algorithm into *potfit* as a new local minimization algorithm. The extended *potfit* was tested by generating a training set using the KIM Environment-Dependent Interatomic Potential (EDIP) model for silicon and using *potfit* to recover the potential parameters from different initial guesses. The results show that EDIP is a "sloppy model" in the sense that its predictions are insensitive to some of its parameters, which makes fitting more difficult. We find that the geodesic LM algorithm is particularly efficient for this case. The extended *potfit* code is the first step in developing a KIM-based fitting framework for interatomic potentials for bulk and two-dimensional materials. The code is available for download via https://www.potfit.net.

‡ Corresponding author.

1. Introduction

An interatomic potential (IP) is a model for approximating the quantum-mechanical interaction of electrons and nuclei in a material through a parameterized functional form that depends only on the positions of the nuclei. IPs such as the Lennard-Jones [1,2] and Morse [3] potentials were initially introduced as part of theoretical studies of material behavior in the early 20th century. Interest in IPs renewed in the 1960s following the development of Monte Carlo and molecular dynamics (MD) simulation methods. This led to the development of IPs for a variety of material systems, a trend that accelerated in the 1980s with the formulation of a large number of increasingly complex IPs. (See [4] for more on the history and scientific development of IPs and MD.) Even complex IPs are far less costly to compute than a first principles approach that involves solving the Schrödinger equation for the quantum system. This enables simulations of large systems (millions, billions and even trillions of atoms depending on the IP) for very short times or moderate-size systems for longer times (tens or hundreds of nanoseconds). Such simulations can tackle problems that are inaccessible to quantum calculations, such as plastic deformation, fracture, atomic diffusion, and phase transformations [5].

Obtaining an accurate and reliable IP constitutes a challenging problem that involves choosing an appropriate functional form and adjusting the parameters to optimally reproduce a relevant training set of experimental and/or quantum data. Traditionally IPs were fitted to reproduce a set of material properties considered important for a given application, such as the cohesive energy, equilibrium lattice constant, and elastic constants of a given crystal phase. However experience has shown that the *transferability* of such IPs (i.e. their ability to accurately predict behavior that they were not fitted to reproduce) can be limited due to the small number of atomic configurations sampled in the training set (although recent work [6] has shown that this approach can be effective in some cases). Further, as the complexity of IPs increases (both in terms of the functional forms and the number of parameters) it can be difficult to obtain a sufficient number of material properties for the training set. This is particularly true for multispecies systems like intermetallic alloys.

To address these difficulties, Ercolessi and Adams [7] proposed the *force-matching method* in which a training set containing material properties is augmented with the forces on atoms obtained by first principles calculations for a set of atomic configurations. These can be configurations associated with important structures or simply random snapshots of the crystal as the atoms oscillate at finite temperature. By fitting to this information the transferability of the potential is likely enhanced since it is exposed to a larger cross-section of configuration space. The issue of insufficient training data is also resolved since as many configurations as needed can be easily generated. This makes it possible to increase the number of parameters and in fact in the original Ercolessi–Adams potential for aluminum [7] the functional forms were taken to be cubic splines with the spline knots serving as parameters. This gives maximum freedom to the fitting process (at the expense of a clear physical motivation for the functional form). The force-matching method has been widely used since its introduction and there are a number of open source implementations available. The *ForceFit* [8] and *ForceBalance* [9] packages target organic-chemistry applications, while *potfit* [10–12] focuses on solid-state physics.

IP fitting programs (such as those mentioned above) define a "cost function" quantifying the difference between the training set data and IP predictions and use global and/or local minimization algorithms to reduce the cost as much as possible. For example, the *potfit* program uses simulated annealing for global minimization followed by a local polish using a conjugate gradient method. A difficulty associated with this procedure is that IPs are nonlinear

functions that are often "sloppy" in the sense that their predictions are insensitive to certain combinations of their parameters [13]. These soft modes in parameter space can cause the minimization algorithms to fail to converge. Recently an understanding of sloppy models based on ideas from differential geometry has led to efficient methods for fitting such models [14] The basic idea is that the parameters of a model (like an IP) define a manifold in the space of data predictions. Fitting the IP then corresponds to finding the point on the manifold closest to the training set. Using these ideas Transtrum et al. [14] augmented the Levenberg–Marquardt algorithm with a geodesic acceleration adjustment to improve convergence. The new geodesic Levenberg–Marquardt (LM) algorithm is likely to be more efficient and more likely to converge for sloppy model systems than conventional approaches. Geodesic LM has been applied to a variety of applications in physics and biology, but until now not to the IP fitting problem.

Once an IP is developed it must be integrated into an MD code (or other atomic-scale simulation package). Traditionally this was done on a code-by-code basis. This creates a barrier for the universal usability of IPs. Users are typically limited by the IPs that happen to be available in the package that they choose. The same applies to IP fitting programs, which support a limited set of functional forms that they can fit. Extending codes to support new IPs is a time consuming and error prone process. It also leads to an inevitable uncontrolled mutation of IPs as modifications to functional forms and/or parameters are introduced (often without adequate documentation). More recently, emphasis have been placed on archival storage of IPs within projects such as the NIST Interatomic Potential Repository (IPR) [15] and the Open Knowledgebase of Interatomic Models (KIM) project [16–18]. The NIST IPR stores tabulated parameters for a limited number of functional forms (mostly the embedded-atom method (EAM) [19–21]). OpenKIM hosted at `https://openkim.org` is a more general effort in which the IP implementation itself is stored in a format that conforms to an application programming interface (API) standard developed as part of the KIM project [22]. The KIM API enables any IP stored in KIM to work seamlessly with any KIM-compliant simulation code including ASAP [23], DL_POLY [24, 25], GULP [26, 27], IMD [28–30], LAMMPS [31, 32], libAtoms+QUIP [33] and Quasicontinuum [34]. In addition IPs stored in OpenKIM are exhaustively tested by computing their prediction for a large number of material properties. These calculations are performed by "KIM Tests" also archived in OpenKIM. Users can extend the OpenKIM system by uploading their own Tests. The predictions of the IPs can be explored using text-based searches and a user-extendible visualization system [35].

In this paper we describe a KIM-compliant implementation of the *potfit* program designed for fitting sloppy IPs. The work has two main objectives. First, we extend the *potfit* program to support the KIM API making it possible to fit IPs with arbitrary functional forms that are portable to a large number of KIM-compliant simulation codes. Second, we explore the efficiency of the geodesic LM algorithm for fitting IPs using the force-matching method within the *potfit* framework. We test the new implementation on a KIM-compliant implementation of the environment-dependent interatomic potential (EDIP) for silicon [36–38]. We use the EDIP model to generate a training set of force configurations and explore the ability of *potfit* to converge back to the original EDIP parameters from different initial guesses. We find that EDIP has the properties of a sloppy model and that the geodesic LM algorithm can converge to the correct results in situations where other minimizers in *potfit* fail.

The paper is organized as follows. In sections 2 and 3, we briefly describe the *potfit* program and the KIM project, respectively. This is followed by a discussion of the modifications that have been made to *potfit* in order to support the KIM API and the geodesic LM algorithm. In section 5, we present the application of the extended *potfit* to the EDIP

model. We conclude in section 6 with a summary and a discussion of future directions.

2. The *potfit* program for fitting interatomic potentials

The *potfit* program [10–12] (`http://www.potfit.net`) is an open source implementation of the force-matching method. It optimizes an IP's parameters by minimizing a cost function defined as the weighted sum of squares of deviations between calculated and reference quantities for forces, energies and/or stresses. (The particular form of the cost function used in this paper is described in section 5.) *potfit* consists of two largely independent components: IP implementations, and least-squares minimization. This program architecture makes it comparatively easy to add either an IP model or a minimization algorithm, which facilitated both aspects of the current work.

potfit supports both analytic and tabulated IPs. An analytic IP has a fixed functional form with certain free parameters. Consider for example the simplest pair potential form where the total energy of n interacting atoms is given by

$$V = \frac{1}{2}\sum_{\alpha=1}^{n}\sum_{\beta=1}^{n}\phi(r_{\alpha\beta}),$$

in which $r_{\alpha\beta}$ is the distance between atoms α and β, and $\phi(r_{\alpha\beta})$ is the energy in the bond connecting them. An example of an analytic function, is the Lennard-Jones potential for which

$$\phi(r) = 4\varepsilon\left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6}\right],$$

where ε and σ are the IP parameters. In contrast, for a tabulated IP, the functional form is arbitrarily defined as an interpolation between sampling points that serve as the IP parameters. For example, for a cubic spline, given a set of distances $(r_1, r_2, \cdots, r_N)$ and a corresponding set of function values $(\phi_1, \phi_2, \cdots, \phi_N)$, the functional form in the interval $[r_i, r_{i+1}]$ is defined as

$$\phi(r) = a_i(r-r_i)^3 + b_i(r-r_i)^2 + c_i(r-r_i) + d_i,$$

where the coefficients (a_i, b_i, c_i, d_i) are obtained from the conditions that $\phi(r_i) = \phi_i$ and $\phi(r_{i+1}) = \phi_{i+1}$ (with different conditions at the ends that define the type of cubic spline). In this scenario the tabulated function values $(\phi_1, \phi_2, \cdots, \phi_N)$ are the IP parameters. The form of interpolation used (cubic, quartic, quintic, or otherwise) can have a strong influence on the predictions of the IP [39]. This is why the interpolation is considered part of the IP by OpenKIM (see section 3) and stored along with the parameters. The *potfit* program supports various functional forms including a variety of pair potentials, embedded atom method (EAM) potentials [19–21], angular-dependent potentials (ADP) [40], bond-order potentials (Tersoff potential [41, 42]), and induced dipole potentials (Tangney-Scandolo [43]).

Once an IP is selected, different local and global minimization algorithms are available within *potfit* to find the optimal set of parameters that minimize the cost function. The *local* minimization algorithm is Powell's method [44], which is a variant of the conjugate gradient (CG) method. Standard CG requires the derivative of the cost function with respect to the IP parameters. The advantage of Powell's method is that it constructs conjugate directions without requiring derivatives. Powell's algorithm begins with an initial guess for the IP parameters and converges to a nearby minimum. For highly nonlinear IPs it is possible that different initial guesses will lead to different solutions. To address this, *global* minimization algorithms attempt to explore the parameter space to find the deepest minimum. *potfit* includes simulated annealing and differential evolution in this class of methods. See the original *potfit* publications [10–12] for detailed descriptions of the optimization algorithms.

3. The Open Knowledgebase of Interatomic Models (KIM) project

The KIM project [16–18] is an international effort aimed at improving the reliability and accuracy of molecular simulations using IPs. A schematic of the OpenKIM cyberinfrastructure is displayed in figure 1. First and foremost OpenKIM provides a repository for IPs at `https://openkim.org`. Within the OpenKIM system IPs are referred to more generically as *models*. A KIM Model consists of the computer implementation of the IP along with any parameters. A KIM IP can be a stand-alone model, or a *model driver* that reads in different parameter files to define different models. All content in the OpenKIM system is archived subject to strict provenance control and can be cited using permanent links. For example, later in this article we will be analyzing the EDIP model for silicon [36–38] archived in OpenKIM [16, 45, 46]. The citations contain a unique 12-digit KIM identifier and 3-digit version number. This makes it possible to access the actual IP used in this publication at any later date to reproduce the calculations — an ability lacking prior to OpenKIM archiving. All content within the OpenKIM repository is citeable in this manner and is accessible to external users through web queries.

In addition to archiving IPs, OpenKIM is tasked with determining the suitability of an IP for a given application. Since IPs are approximate, their transferability to phenomena they were not fitted to reproduce is limited. To help define a rigorous method for IP selection, OpenKIM includes a user-extendible suite of *KIM Tests* for characterizing an IP. Each KIM Test computes the prediction of an IP for a well-defined material property (such as the surface free energy for a given crystal face under specified temperature and pressure). Each IP in the system is mated with all compatible KIM Tests using an automated system called the *Processing Pipeline* [47]. Whenever a new Model or Test is uploaded to OpenKIM by members of the community, the pipeline launches to perform the new calculations. The results are stored in the OpenKIM repository and can be explored and compared with first principles and experimental reference data using text-based and graphical visualization tools [35]. Machine learning based tools are being developed to use the accumulated information on the accuracy of an IPs predictions to help select for applications of choice [48].

The operation of the pipeline involves the coupling of pairs of computer codes, Model and Test, which may be written in different computer languages and need to exchange information. Specifically, the Test repeatedly calls the Model with new input (typically positions and species of atoms) and requires the generated output (typically energy and forces on atoms). This is made possible because all codes in OpenKIM conform to the KIM API [22]. This is a standard developed as part of the KIM project for atomistic simulation codes to exchange information. The KIM API is a lightweight efficient interface where only pointers are exchanged between the interfacing codes, which can be written in any supported language

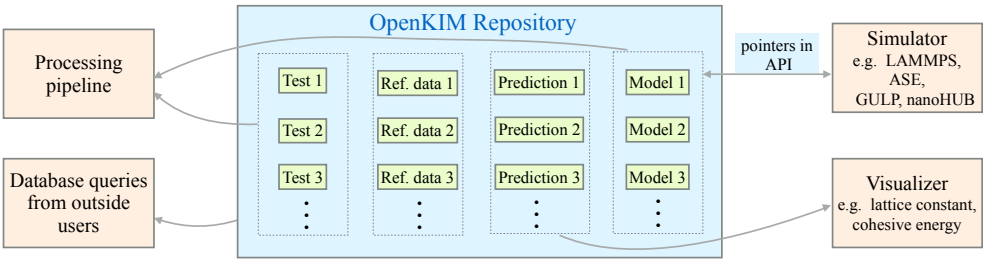

Figure 1: OpenKIM framework

(C, C++, Fortran 77/90/95/2003, Python). A number of major simulation codes currently support the KIM API (listed in section 1). Anyone of these codes can work seamlessly (without having to be recompiled) with any KIM Model stored on the same computer as a dynamically linked library. (This is akin to a CD being placed in a player from any manufacturer and having its content played.) This greatly increases the portability of IPs. It allows researchers to select the appropriate IP for their application and use it with their code of choice, rather than being limited to the IPs that happen to come with it.

4. Extensions to *potfit*: Geodesic LM and KIM-compliance

An extended version of the *potfit* code has been developed that includes a geodesic LM optimizer and support for the KIM API. The new code is available via `https://www.potfit.net`. Details of the extensions are given below.

4.1. Geodesic Levenberg–Marquardt (LM) algorithm

As explained in the introduction, IPs often have the characteristics of sloppy models. This can cause the standard minimization algorithms in *potfit* to fail to converge. As an alternative we have incorporated into *potfit* the newly developed geodesic LM algorithm [14, 49–51]. Geodesic LM is an efficient method offering a higher likelihood of convergence for sloppy models. The performance of the geodesic LM method is discussed in section 5.4. The method is described next.

The least-squares optimization problem is to find the parameters $\boldsymbol{\theta} = \{\theta_1, \ldots, \theta_N\}$ that minimize the cost function $C : \mathbb{R}^N \mapsto \mathbb{R}$:

$$C(\boldsymbol{\theta}) = \frac{1}{2}\sum_{m=1}^{M}(r_m(\boldsymbol{\theta}))^2 = \frac{1}{2}\|\boldsymbol{r}(\boldsymbol{\theta})\|^2 = \frac{1}{2}\boldsymbol{r}(\boldsymbol{\theta})^{\mathrm{T}}\boldsymbol{r}(\boldsymbol{\theta}), \tag{1}$$

where the residual $\boldsymbol{r} : \mathbb{R}^N \mapsto \mathbb{R}^M$ is an M-dimensional vector function of N parameters. The gradient of C is

$$\frac{\partial C}{\partial \theta_i}(\boldsymbol{\theta}) = \sum_{m=1}^{M} r_m \frac{\partial r_m}{\partial \theta_i} \quad \Leftrightarrow \quad \frac{\partial C}{\partial \boldsymbol{\theta}}(\boldsymbol{\theta}) = \boldsymbol{J}^{\mathrm{T}}\boldsymbol{r}, \tag{2}$$

where $\boldsymbol{J}$ is the Jacobian of $\boldsymbol{r}$, and the Hessian of C is,

$$\frac{\partial^2 C}{\partial \theta_i \theta_j}(\boldsymbol{\theta}) = \sum_{m=1}^{M}\left(\frac{\partial r_m}{\partial \theta_i}\frac{\partial r_m}{\partial \theta_j} + r_m\frac{\partial^2 r_m}{\partial \theta_i \theta_j}\right) \quad \Leftrightarrow \quad \frac{\partial^2 C}{\partial \boldsymbol{\theta}\partial\boldsymbol{\theta}}(\boldsymbol{\theta}) = \boldsymbol{J}^{\mathrm{T}}\boldsymbol{J} + \boldsymbol{r}^{T}\boldsymbol{K}, \tag{3}$$

where $\boldsymbol{K}$ is the second derivative of $\boldsymbol{r}$ with respect to $\boldsymbol{\theta}$ (a third-order matrix).

Before introducing geodesic LM, let us consider the more basic Gauss–Newton method. This approach is based on a linear approximation of $\boldsymbol{r}$ [52],

$$\boldsymbol{r}(\boldsymbol{\theta} + \delta\boldsymbol{\theta}) \approx \boldsymbol{r}(\boldsymbol{\theta}) + \boldsymbol{J}(\boldsymbol{\theta})\delta\boldsymbol{\theta}. \tag{4}$$

Substituting (4) into (1) we see that,

$$\begin{aligned} C(\boldsymbol{\theta} + \delta\boldsymbol{\theta}) &= \frac{1}{2}\boldsymbol{r}(\boldsymbol{\theta} + \delta\boldsymbol{\theta})^{\mathrm{T}}\boldsymbol{r}(\boldsymbol{\theta} + \delta\boldsymbol{\theta}) \\ &= \frac{1}{2}(\boldsymbol{r} + \boldsymbol{J}\delta\boldsymbol{\theta})^{\mathrm{T}}(\boldsymbol{r} + \boldsymbol{J}\delta\boldsymbol{\theta}) \\ &= C(\boldsymbol{\theta}) + \delta\boldsymbol{\theta}^{\mathrm{T}}\boldsymbol{J}^{\mathrm{T}}\boldsymbol{r} + \frac{1}{2}\delta\boldsymbol{\theta}^{\mathrm{T}}\boldsymbol{J}^{\mathrm{T}}\boldsymbol{J}\delta\boldsymbol{\theta}. \end{aligned} \tag{5}$$

Taking the first derivative of (5) with respect to $\delta\boldsymbol{\theta}$ and setting it to zero, we have,

$$\delta\boldsymbol{\theta} = -(\boldsymbol{J}^{\mathrm{T}}\boldsymbol{J})^{-1}\boldsymbol{J}^{\mathrm{T}}\boldsymbol{r}. \tag{6}$$

In (6) $\delta\boldsymbol{\theta}$ is a local decent direction of C, thus we can obtain the local minimum of C by solving (6) iteratively. This is the Gauss–Newton method.

Levenberg [53] suggested a "damped version" of (6) where the parameters are iteratively updated according to

$$\delta\boldsymbol{\theta} = -(\boldsymbol{J}^{\mathrm{T}}\boldsymbol{J} + \lambda\boldsymbol{D})^{-1}\boldsymbol{J}^{\mathrm{T}}\boldsymbol{r}, \quad \lambda \geq 0. \tag{7}$$

Here λ is a (non-negative) damping parameter, and $\boldsymbol{D} = \boldsymbol{I}$ is the identity. The Levenberg method is an interpolation between the Gauss–Newton algorithm and the steepest decent algorithm. For small λ, (7) reduces to the Gauss–Newton equation in (6); whereas for large λ, $\delta\boldsymbol{\theta} \approx -\lambda^{-1}\boldsymbol{J}^{\mathrm{T}}\boldsymbol{r}$, which lies along the gradient (i.e. the steepest descent direction) in (2).

In order to overcome slow convergence in directions of small gradients, Marquardt [54] proposed taking $\boldsymbol{D} = \mathrm{diag}(\boldsymbol{J}^{\mathrm{T}}\boldsymbol{J})$, where $\mathrm{diag}(\boldsymbol{A})$ returns a diagonal matrix with elements equal to the diagonal elements of the square matrix $\boldsymbol{A}$). The downside of this approach is that it greatly increases the susceptibility for parameter evaporation [50], i.e. the algorithm pushes the parameters to infinite values without finding a good fit around saddle points.

To improve the efficiency of the Levenberg–Marquardt (LM) method, Transtrum and Sethna [14, 49–51] proposed the geodesic LM algorithm in which the minimization step size is constrained and the residual is modified to include a second-order (geodesic acceleration) correction,

$$\boldsymbol{r}(\boldsymbol{\theta} + \delta\boldsymbol{\theta}) \approx \boldsymbol{r}(\boldsymbol{\theta}) + \boldsymbol{J}(\boldsymbol{\theta})\delta\boldsymbol{\theta} + \frac{1}{2}\delta\boldsymbol{\theta}^{\mathrm{T}}\boldsymbol{K}(\boldsymbol{\theta})\delta\boldsymbol{\theta}. \tag{8}$$

The expansion of the cost function is then

$$C(\boldsymbol{\theta} + \delta\boldsymbol{\theta}) = \frac{1}{2}\boldsymbol{r}(\boldsymbol{\theta} + \delta\boldsymbol{\theta})^{\mathrm{T}}\boldsymbol{r}(\boldsymbol{\theta} + \delta\boldsymbol{\theta}) + \frac{1}{2}\lambda\delta\boldsymbol{\theta}^{\mathrm{T}}\boldsymbol{D}\delta\boldsymbol{\theta}, \tag{9}$$

where a Lagrange multiplier enforces the constraint on the step size $\delta\boldsymbol{\theta}$. Taking the first derivative of (9) with respect to $\delta\boldsymbol{\theta}$ and setting it to zero, we have

$$J_{mi}r_m + (J_{mi}J_{mj} + r_m K_{mij} + \lambda D_{ij})\delta\theta_j + (J_{mk}K_{mij} + \frac{1}{2}J_{mi}K_{mkj})\delta\theta_j\delta\theta_k = 0, \tag{10}$$

where the indices are included explicitly to avoid ambiguity and Einstein's summation convention over repeated indices is adopted. The step size $\delta\theta_i$ is written as a sum of two parts:

$$\delta\theta_i = \delta\theta_i^{(1)} + \delta\theta_i^{(2)}, \tag{11}$$

where $\delta\theta_i^{(1)}$ is given by (7) and $\delta\theta_i^{(2)}$ includes the remaining terms. Substituting (11) and (7) into (10), $\delta\theta_i^{(2)}$ can be solved as,

$$\begin{aligned}\delta\theta_i^{(2)} &= -\frac{1}{2}(J_{mi}J_{mj} + \lambda D_{ij})^{-1}J_{pj}K_{pkl}\delta\theta_k^{(1)}\delta\theta_l^{(1)} \\ &= -\frac{1}{2}(J_{mi}J_{mj} + \lambda D_{ij})^{-1}r_p'',\end{aligned} \tag{12}$$

where some small terms are ignored (see [50] for details) and the directional second derivative $r_p'' = K_{pkl}\delta\theta_k^{(1)}\delta\theta_l^{(1)}$ is used.

As seen from (12), the geodesic acceleration correction only depends on the directional second derivative oriented along the first order correction $\delta\boldsymbol{\theta}^{(1)}$. This feature is very important

because the cost to compute the directional second derivative is reasonably small and hence will not overly add to the computational burden.

Finally, in order to prevent parameter evaporation in the geodesic LM algorithm and to increase the likelihood of convergence, allowable step sizes must satisfy:

$$\frac{2\|\delta\boldsymbol{\theta}^{(2)}\|}{\|\delta\boldsymbol{\theta}^{(1)}\|} \leq \alpha, \tag{13}$$

where α is some parameter (usually smaller than 1) and whose optimal value depends on the specific problem.

An open source implementation of the geodesic LM algorithm has been made available by Transtrum [55]. In addition to the geodesic acceleration correction, the geodesic LM package includes different options for updates to the damping parameter and the damping matrix, different conditions for accepting uphill steps (i.e. steps that increase C), and whether or not to use Broyden's method [56] to lower the Jacobian update frequency. In this work, the damping matrix is set to $\boldsymbol{I}$, Umrigar and Nightingale's method [50] is used to update the damping parameter and to decide whether to accept uphill steps, and we do not employ Broyden's method to update the Jacobian because the computational cost in this problem is not too high.

4.2. KIM-compliance

The original version of *potfit* is limited to using IPs built into the program. The philosophy underlying the KIM API is different. It is based on the idea of separation between the simulation code (called a "simulator" in KIM parlance) and the IP (called the KIM Model). Whenever the simulator (*potfit* in this case) requires the energy or forces for an atomic configuration, it calls a KIM Model with the necessary input (e.g. coordinates and species of the atoms) and receives back the required output (e.g. the forces on the atoms). Enforcing this separation makes it possible for the simulator to work with any compatible KIM Model in plug-and-play fashion.§ The KIM API is the standard for the exchange of information between the communicating codes. In practice this works by having all data exchanged between the simulator and IP stored in a data structure called the "KIM API Object." The simulator and IP can access and change this data using KIM API library routines. The only direct data transfer required between the programs is for the simulator to pass to the KIM Model the pointer to the location in memory of the KIM API Object. This makes the KIM API extremely efficient and provides it with cross language compatibility (as described in section 3). This means that *potfit*, which is written in C, can work efficiently and seamlessly with KIM Models written in C, C++, Fortran 77/90/95/2003, and Python.

The *potfit* program has an unusual requirement for a simulator in that it needs to modify the parameters of the IP that it is optimizing. This is supported by the KIM API through the ability of KIM Models to "publish their parameters." A KIM Model that does so has in its KIM descriptor file a list of its parameters identified by name with the prefix "PARAM_FREE_". The *potfit* user can select any of these parameters to optimize. This is specified in the input file to *potfit* along with the initial guesses for the parameters, and the KIM Model identifier (see section 3).

The operation of the KIM-compliant implementation of *potfit* is described in figure 2. The program reads in the input defining the optimization problem and then generates one or

§ A KIM Model is compatible with a simulator if it supports the required species and boundary conditions that the simulator requires. This is automatically determined by the KIM API using metadata in "KIM descriptor" files that accompany the simulator and KIM Model.

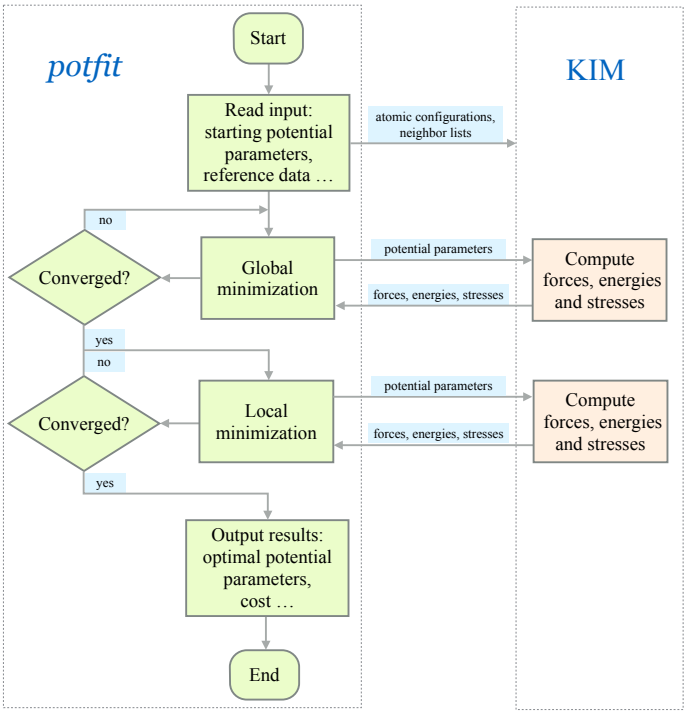

Figure 2: Flowchart of the KIM-compliant *potfit* program

more KIM API Objects containing the training set to be used for the optimization. At each step in the optimization whenever the IP parameters need to be changed, the following steps are performed using KIM API library calls:

(i) The KIM Model parameters are changed in the KIM API Object;

(ii) The KIM Model is called to compute the new energy and forces;

(iii) *potfit* acquires the energy and forces from the KIM API Object.

When the optimization is complete, *potfit* outputs the final set of parameters and writes out a new KIM Model, which can be uploaded to `https://openkim.org` if desired.

5. Application of *KIM-potfit* to EDIP

In this section we test the performance of *KIM-potfit* by applying it to the optimization of the EDIP model for silicon [36–38] archived in OpenKIM [16, 45, 46]. The training set consists of the energy and forces computed with the original EDIP model for a periodic configuration of $5 \times 5 \times 5$ conventional unit cells for a total of 1000 atoms perturbed off the ideal diamond structure positions. Our objective is to see whether starting from different initial guesses for the parameters, the optimization procedure is able to recover the EDIP parameters.

5.1. Cost function

As explained in section 2, IP optimization corresponds to minimizing a cost function quantifying the difference between the IP's predictions and a training set of reference data.

Our training set consists of a single atomic configuration with $n = 1000$ atoms. The cost function used in our analysis is

$$C(\boldsymbol{\theta}) = \frac{1}{2} w_1 \sum_{\alpha=1}^{n} \|\boldsymbol{f}_\alpha - \boldsymbol{f}_\alpha^0\|^2 + \frac{1}{2} w_2 (E - E^0)^2, \tag{14}$$

where $\boldsymbol{f}_\alpha = \boldsymbol{f}_\alpha(\boldsymbol{\theta})$ and $\boldsymbol{f}_\alpha^0$ are the predicted and reference force on atom α, $E = E(\boldsymbol{\theta})$ and E^0 are the predicted and reference energy, and w_1 and w_2 are weights. In our example we take $w_1 = 1\,(\text{Å/eV})^2$, $w_2 = 1\,(1/\text{eV})^2$ to give a unitless cost function. The reference forces and energy are defined as $\boldsymbol{f}_\alpha^0 = \boldsymbol{f}_\alpha(\boldsymbol{\theta}^0)$ and $E^0 = E(\boldsymbol{\theta}^0)$, where $\boldsymbol{\theta}^0$ are the original EDIP parameters. EDIP's functional form and parameters along with the training set used are given in the supplementary information accompanying this article.

The Hessian of the cost function at the original parameters $\boldsymbol{\theta}^0$ follows from (3) as

$$H_{ij}^0 = \left.\frac{\partial^2 C}{\partial\theta_i \partial\theta_j}\right|_{\boldsymbol{\theta}^0} = w_1 \sum_{\alpha=1}^{M} \left.\frac{\partial \boldsymbol{f}_\alpha}{\partial\theta_i}\right|_{\boldsymbol{\theta}^0} \cdot \left.\frac{\partial \boldsymbol{f}_\alpha}{\partial\theta_j}\right|_{\boldsymbol{\theta}^0} + w_2 \left.\frac{\partial E}{\partial\theta_i}\right|_{\boldsymbol{\theta}^0} \left.\frac{\partial E}{\partial\theta_j}\right|_{\boldsymbol{\theta}^0}, \tag{15}$$

where we have used $(\boldsymbol{f}_\alpha - \boldsymbol{f}_\alpha^0)\big|_{\boldsymbol{\theta}^0} = 0$ and $(E - E^0)\big|_{\boldsymbol{\theta}^0} = 0$. We will use the Hessian in the next section to evaluate the sensitivity of EDIP predictions to its parameters.

5.2. EDIP sensitivity analysis

It is a common feature of IPs that the prediction of a model (including the cost function) are weakly dependent on certain directions in parameter space, i.e. the models are *sloppy* [57]. This is demonstrated schematically in figure 3 showing the contour plot of a cost function with only two parameters. Clearly the function varies more slowly along the $\boldsymbol{\Lambda}_2$ eigendirection than along $\boldsymbol{\Lambda}_1$. As a result of this structure of $C(\boldsymbol{\theta})$, convergence along certain directions can be incomplete. For example in the figure, the minimum is located at O but a convergence could terminate at a set of parameters offset to the optimal ones by $\Delta\boldsymbol{\theta}$. An indication of whether this is occurring can be obtained by computing the angle between $\Delta\boldsymbol{\theta}$ and the eigenvectors of the Hessian $\partial^2 C/\partial\theta_i\partial\theta_j$ evaluated at O,

$$\cos\alpha_i = \frac{\Delta\boldsymbol{\theta}\cdot\boldsymbol{\Lambda}_i}{\|\Delta\boldsymbol{\theta}\|}. \tag{16}$$

Point A and point B (corresponding to another fit) are equally close to the eigendirections because they form the same angles.

In general, the eigendirections associated with smaller eigenvalues are more "sloppy" and we would expect larger scatter in the $\cos\alpha_i$ values computed in (16) for an ensemble of fits. Thus qualitatively we expect a behavior similar to that shown in figure 4, where the scatter in $\cos\alpha_i$ obtained from a large number of fitting attempts with different initial guesses would occupy a region similar to Region I in the figure with larger scatter along directions associated with smaller eigenvalues. This would be the signature of a sloppy model.

We perform this analysis for EDIP. EDIP has 13 parameters, but only 11 of them are used in this study with the two cutoffs excluded. The training set is described in section 5.1. We take an ensemble of 100 initial guesses obtained by adding to the original EDIP parameters random numbers ρ drawn from a Gaussian distribution with average zero and a standard deviation of 0.15, i.e. $\theta_i := \theta_i^0 \times (1+\rho), i = 1,\ldots,11$, and perform the fitting procedure for each one. The relative errors $\Delta\theta_i/\theta_i^0 = (\theta_i^* - \theta_i^0)/\theta_i^0$ (for $i = 1,\ldots,11$) between the parameters $\boldsymbol{\theta}^*$ obtained in the fitting procedure and the original EDIP parameters $\boldsymbol{\theta}^0$ for all 100 optimizations are shown in figure 5a. No clear trend is apparent. Instead let us plot the cosine defined in (16) as done schematically in figure 4. To improve the scaling of the plot let

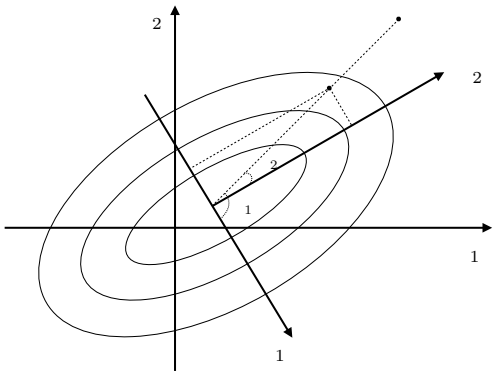

Figure 3: Contour plot of a cost function with two parameters. Λ_1 and Λ_2 are the eigendirections of the Hessian of the cost function corresponding to the larger and smaller eigenvalues, respectively. Points *A* and *B* represent two fits.

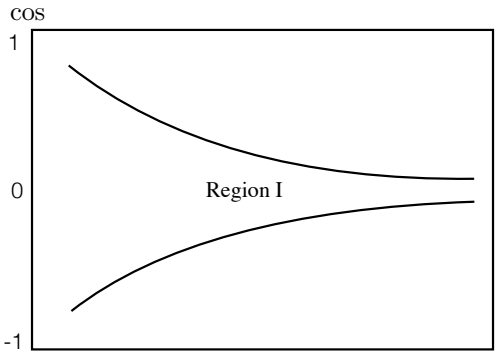

Figure 4: Expected region for normalized projected parameter errors if the fits are aligned along the sloppy direction(s). The horizontal axis corresponds to the eigenvalues of the Hessian sorted in increasing order.

us modify the definition to use the logarithm of the parameters, $\hat{\theta}_i = \log \theta_i$. (This is important when different parameters have very different magnitudes.) The definition equivalent to (16) is then

$$\cos \hat{\alpha}_i = \frac{\Delta\hat{\boldsymbol{\theta}} \cdot \hat{\boldsymbol{\Lambda}}_i}{\|\Delta\hat{\boldsymbol{\theta}}\|}, \tag{17}$$

where $\Delta\hat{\boldsymbol{\theta}} = \hat{\boldsymbol{\theta}}^* - \hat{\boldsymbol{\theta}}^0$ and $\hat{\boldsymbol{\Lambda}}_{\boldsymbol{i}}$ is the *i*th eigendirection of the Hessian of the cost function C in the logarithm parameter space, $\partial^2 C/\partial\hat{\theta}_i\partial\hat{\theta}_j = \theta_i\theta_j\partial^2 C/\partial\theta_i\partial\theta_j$. The results for EDIP are plotted in figure 5b, and are very much in line with the expected behavior in figure 4. In particular, we see that the spread is mostly accounted for by the first four eigenterms each of which corresponds to a particular combination of parameters as quantified by the eigendirections.

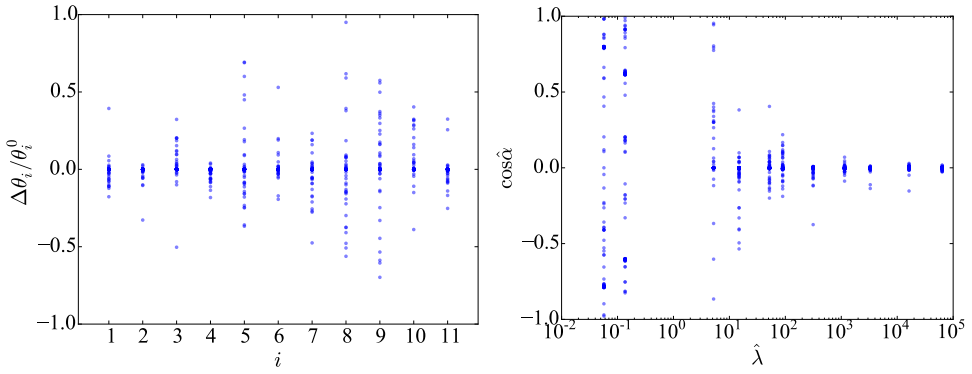

(a) (b)

Figure 5: Errors in fitting the EDIP parameters. (a) Relative errors. (b) Normalized error projected on the eigendirections of the Hessian of the cost function in the logarithm parameter space.

5.3. Cost along eigendirections

The results in the previous section support the conjecture that the observed scatter in results is due to sloppy directions. However, since the Hessian is a local measure of the structure of C near the optimal solution, it is possible that there are other minima further away that can trap the optimization process. To explore this possibility, we plot the cost function along the different eigendirections,

$$\widehat{C}_i(s) = C(\boldsymbol{\theta}^0 + s\boldsymbol{\Lambda}_i), \tag{18}$$

where as before $\boldsymbol{\theta}^0$ corresponds to the original EDIP parameters and s is a scalar factor that can be positive or negative. The results are plotted in figure 6 for all eigendirections from the smallest eigenvalue 1 to the largest eigenvalue 11. The cost $\widehat{C}$ varies more slowly along the eigendirections associated with the five smallest eigenvalues than along other directions, i.e. the cost function is more flat in these five directions. (Note that the scale of the y-axis is not the same in the plots, and the first five plots have a significantly smaller range than the others.) This is in agreement with the previous observation that the spread of the fitted parameters is larger along these directions. We also see that an additional local minimum (aside from the one associated with original parameters at $s = 0$) is only apparent along the positive 11 direction at a significantly higher cost function.‖ This minimum can trap the optimization algorithm for initial guesses in its vicinity, however over a large range near the original parameters there is only a single minimum. This suggests that a local minimizer is sufficient for this IP (using different initial guesses), which can reduce the fitting time drastically compared with global minimization efforts.

5.4. Performance of minimization algorithms

Having established in the previous section that a local minimization is sufficient for EDIP (at least in the vicinity of the original parameters), we explore the efficiency of Powell's method, the geodesic LM algorithm added to *potfit*, and the standard LM algorithm described in [58]

‖ We verified that this is a local minimum by checking the eigenvalues of the Hessian at the minimum and finding that they are all positive.

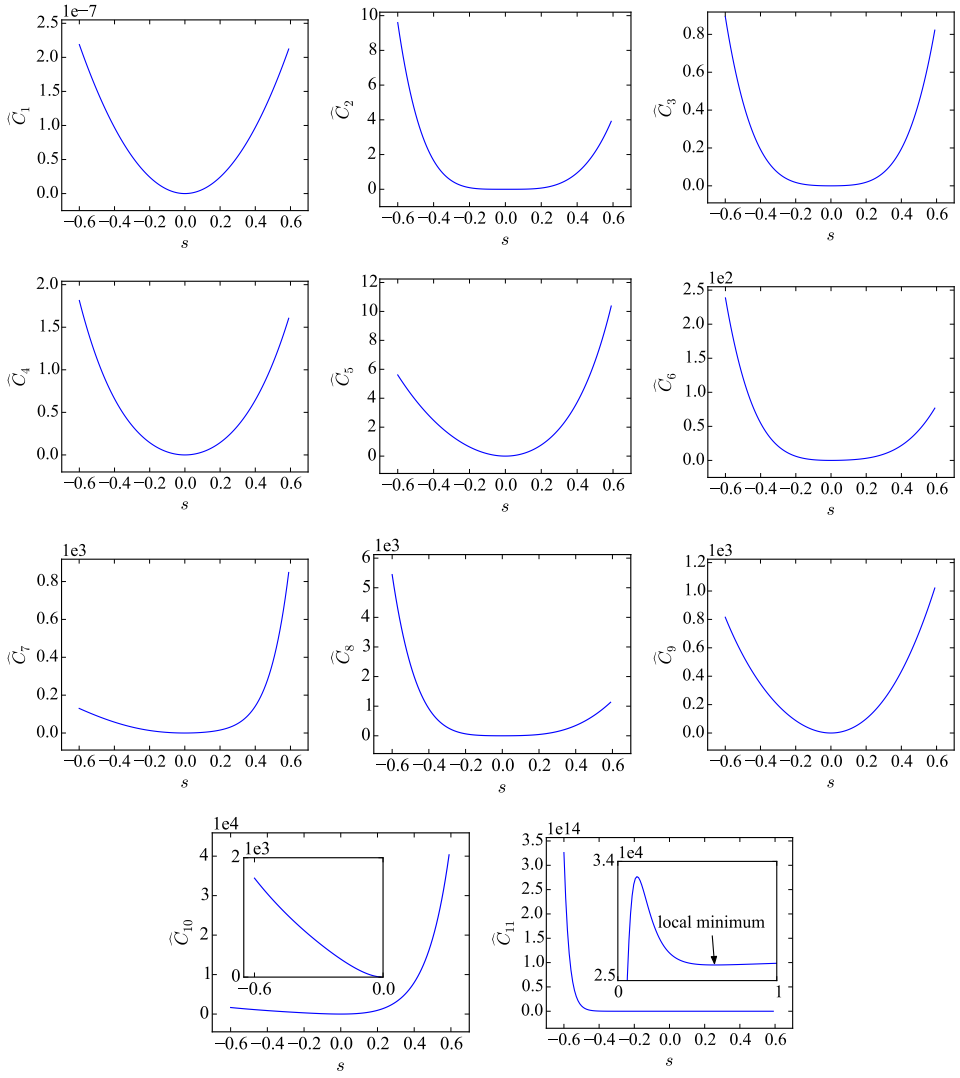

Figure 6: Cost along eigendirections. $\widehat{C}_1$ is associated with the smallest eigenvalue, and $\widehat{C}_{11}$ is associated with the largest. For $\widehat{C}_{10}$ and $\widehat{C}_{11}$ the insets show a zoomed view of the flat portions of the curves.

(see section 4.1 for details of LM and geodesic LM). The three minimizers are applied to fit EDIP. The initial guesses for the parameters are obtained in the same way as discussed in section 5.2. Three sets of 100 initial guesses are constructed by perturbing the original EDIP parameters by random numbers drawn from a Gaussian distribution with zero average and standard deviations of 0.1, 0.2 and 0.3. The results are given in table 1 showing for each minimizer and perturbation magnitude the number of attempts (out of 100) that converged below the cost level indicated on the left. For example for Powell's method with a perturbation magnitude of 0.2, the algorithm converged to a cost below 10^{-3} for 46 out of the 100 initial guesses. We see that geodesic LM outperforms the other two algorithms for all perturbation

Table 1: Statistics on the optimization of EDIP. The number of minimization attempts (out of a 100) that converge to a cost function below the tolerance specified on the left for the three minimizers studied and the three perturbation amplitudes (0.1, 0.2 and 0.3) used to generate the initial guesses.

	Powell			Standard LM			Geodesic LM		
Cost(C)	0.1	0.2	0.3	0.1	0.2	0.3	0.1	0.2	0.3
$C < 10^{-7}$	85	37	7	68	46	9	100	60	14
$C < 10^{-5}$	85	39	8	76	49	9	100	60	14
$C < 10^{-3}$	89	47	9	100	58	16	100	64	21
$C < 10^{-1}$	93	68	23	100	79	25	100	84	33
$C < 10^{0}$	98	78	36	100	85	39	100	90	51
$C < 10^{1}$	100	94	65	100	96	54	100	98	76

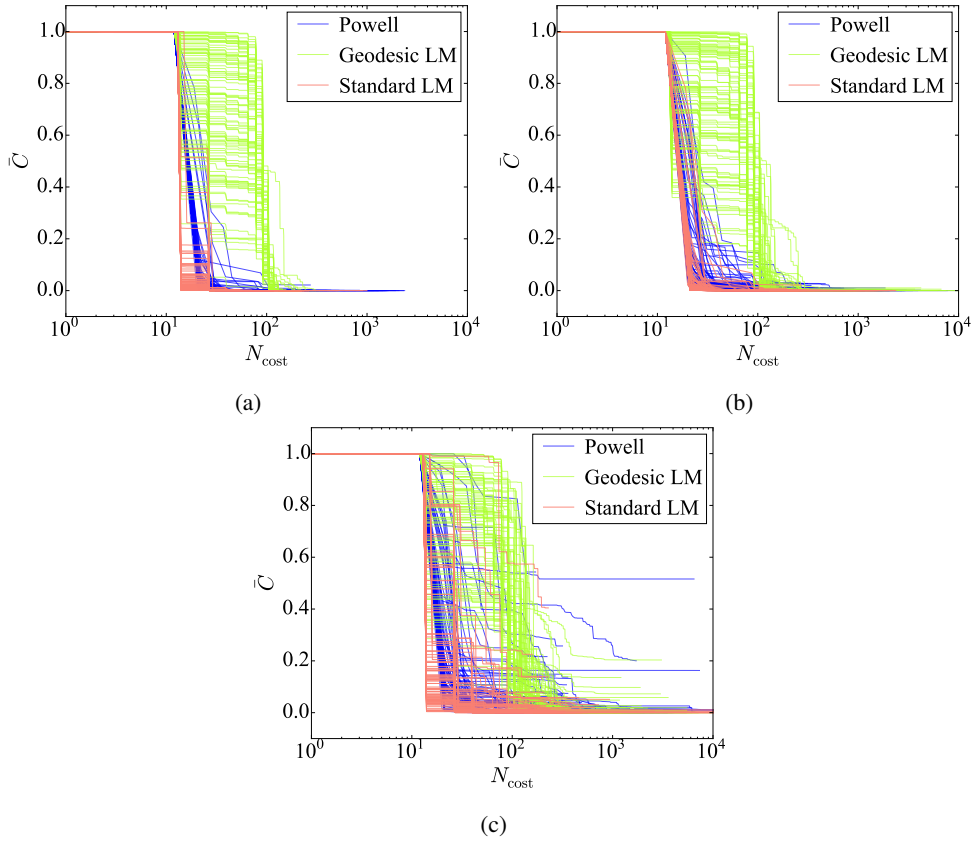

Figure 7: Normalized cost $\bar{C}$ versus the number of cost function evaluations N_{cost}. The initial guesses for the parameters are perturbations off the original EDIP parameters drawn from a Gaussian distribution with zero average and standard deviations of (a) 0.1, (b) 0.2, and (c) 0.3.

amplitudes. In all cases it is able to converge from more initial guess to below the specified tolerance. The standard LM algorithm is superior to Powell in all cases except for the smallest perturbation.

We are also interested in the computational cost of the different minimization methods. The Powell algorithm requires only evaluations of the cost function, whereas LM and geodesic LM require the gradient of the cost function with respect to the parameters. The gradient is computed using numerical differentiation, which only requires evaluations of the cost function. Therefore the number of cost function evaluations (N_{cost}) performed during the minimization is a good metric for comparing the performance of all three methods. Figure 7 shows for each perturbation amplitude, the normalized cost function $\bar{C} = C/C_{\text{init}}$ (where C_{init} is the cost at the start of the minimization) versus N_{cost} for the three methods. The initial flat region is associated with the setup stage of the methods (e.g. building an initial estimate for the Hessian in Powell's method). Examining the plots, we see that Powell and LM are comparable in performance, whereas geodesic LM is about an order of magnitude slower. Thus geodesic LM's superior convergence properties come at the expense of a significant increase in the computational burden. For the current problem this tradeoff is acceptable since geodesic LM converges in most cases in less than 1000 cost function evaluations, which is quite reasonable for the current problem. However, when computation of the cost function is expensive (due to a larger training set or more complex IP) or when more steps are needed to converge as would be expected for IPs with more parameters, the computational cost of geodesic LM may be an issue.

6. Summary and future directions

The accuracy of an IP is critical to atomic-scale simulations. Designing a high-quality transferable IP is a very challenging task due to the complexity of the quantum mechanical behavior that the IP is attempting to mimic. The developer of an IP must select a functional form with adjustable parameters and a suitably-weighted training set of first principles and experimental reference data. The IP is then optimized with respect to its parameters to reproduce the training set as closely as possible in terms of specified cost function. It is then validated against an independent test set of reference data. If the validation results are inadequate, the functional form and/or the training set must be amended and the process repeated until a satisfactory IP is obtained. This process is made more difficult by the fact that many IPs are sloppy models in the sense that their predictions are weakly dependent on certain combinations of their parameters. This can cause optimization algorithms to fail to converge.

potfit is an IP fitting program based on the force-matching method. The idea is to include in the training set the forces on atoms computed from first principles for a set of atomic configurations of interest. This allows the developer to train the IP with a large amount of reference data spanning a larger portion of configuration space. We have extended *potfit* in two ways:

(i) We have incorporated a new optimization method based on the geodesic LM algorithm. Although slower than some other local minimizers it has improved convergence properties for sloppy models. This is the first application of geodesic LM to IP optimization.

(ii) We have adapted *potfit* to be compatible with the KIM API. This allows *potfit* to optimize KIM Models stored in the KIM Repository at `https://openkim.org`. This greatly expands the functional forms that *potfit* can optimize. The resulting IPs are portable and can be used seamlessly with many major software packages that support the KIM API.

We study the effectiveness of the new code by using it to optimize the EDIP model for silicon. This IP is not available within the original *potfit* code, but is available as a KIM Model.

We use EDIP to generate a training set of atomic configurations and test whether *potfit* can reproduce the EDIP parameters from different initial guesses. We make several observations:

(i) EDIP displays the characteristics of a sloppy model. For different initial guesses we find a large scatter in the parameter values. An analysis shows that the scatter is maximal along the eigendirections of the Hessian of the cost function associated with the smallest eigenvalues. Thus the IP's behavior depends primarily on a few stiff combinations of parameters.

(ii) At least in the vicinity of the original parameters the cost function is largely convex with only two local minima. This suggests that a local minimization method with different random initial guesses is sufficient.

(iii) The geodesic LM algorithm was on average twice as likely to converge to the correct solution from different initial guesses than standard LM or Powell's method that tended to get trapped along sloppy directions. However, the improved performance comes at the cost of about an order of magnitude increase in computational expense for geodesic LM. For the current problem the computations were not prohibitive and geodesic LM was the preferred method.

The extended *potfit* code is the first step in developing an IP fitting framework based on the OpenKIM project. Such a framework would make it possible for the developer to draw upon the large selection of KIM Tests (standardized codes for computing material properties) and reference data in OpenKIM. Models developed using this framework *and* the training set used in the fit will then be archived at `https://openkim.org`. The latter is an important advantage as it means that other developers will be able to draw upon previous work and adapt a Model to a new application starting from the training set used in its development. This framework is currently under development for use in fitting IPs for 2D materials [59].

Acknowledgments

We thank Mark Transtrum for his open source implementation of the geodesic LM algorithm, and Colin Clement for useful discussions. The authors also thank Daniel Schopf for his assistance in merging the OpenKIM extension into the main *potfit* distribution. The research was partly supported through the National Science Foundation (NSF) under grants No. PHY-0941493 and DMR-1408211. The authors acknowledge the support of the Army Research Office (W911NF-14-1-0247) under the MURI program.

References

[1] Jones J E 1924 On the determination of molecular fields. I. From the variation of the viscosity of a gas with temperature *Proc. Roy. Soc. A* **106** 441–62

[2] Jones J E 1924 On the determination of molecular fields. II. From the equation of state of a gas *Proc. Roy. Soc. A* **106** 463–77

[3] Morse P M 1929 Diatomic molecules according to the wave mechanics. II. vibrational levels *Phys. Rev.* **34** 57–64

[4] Tadmor E B and Miller R E 2011 *Modeling materials: continuum, atomistic and multiscale techniques* Cambridge University Press, Cambridge

[5] Mishin Y, Farkas D, Mehl M J and Papaconstantopoulos D A 1999 Interatomic potentials for monoatomic metals from experimental data and ab initio calculations *Phys. Rev. B* **59** 3393–407

[6] Nichol A and Ackland G J 2016 Property trends in simple metals: An empirical potential approach *Phys. Rev. B* **93** 184101

[7] Ercolessi F and Adams J B 1994 Interatomic potentials from first-principles calculations: the force-matching method *Europhys. Lett.* **26** 583–8

[8] Waldher B, Kuta J, Chen S, Henson N and Clark A E 2010 Forcefit: a code to fit classical force fields to quantum mechanical potential energy surfaces *J. Comput. Chem.* **31** 2307–16

[9] Wang L P, Chen J and Van Voorhis T 2012 Systematic parametrization of polarizable force fields from quantum chemistry data *J. Chem. Theory Comput.* **9** 452–60

[10] Brommer P and Gähler F 2006 Effective potentials for quasicrystals from ab-initio data *Philos. Mag.* **86** 753–8

[11] Brommer P and Gähler F 2007 Potfit: effective potentials from ab initio data *Model. Simul. Mater. Sci. Eng.* **15** 295

[12] Brommer P, Kiselev A, Schopf D, Beck P, Roth J and Trebin H R 2015 Classical interaction potentials for diverse materials from ab initio data: a review of potfit *Model. Simul. Mater. Sci. Eng.* **23** 074002

[13] Waterfall J J, Casey F P, Gutenkunst R N, Brown K S, Myers C R, Brouwer P W, Elser V and Sethna J P 2006 Sloppy model universality class and the Vandermonde matrix *Phys. Rev. Lett.* **97** 150601

[14] Transtrum M K, Machta B B and Sethna J P 2011 Geometry of nonlinear least squares with applications to sloppy models and optimization *Phys. Rev. E* **83** 036701

[15] Becker C A 2015 Nist interatomic potentials repository project `http://www.ctcms.nist.gov/potentials/`

[16] Tadmor E B, Elliott R S, Sethna J P, Miller R E and Becker C A 2011 The potential of atomistic simulations and the knowledgebase of interatomic models *JOM* **63** 17

[17] Tadmor E B, Elliott R S, Phillpot S R and Sinnott S B 2013 Nsf cyberinfrastructures: A new paradigm for advancing materials simulation *Curr. Opin. Solid State Mater. Sci.* **17** 298–304

[18] 2016 Open knowledgebase of interatomic models (kim) website `https://openkim.org`

[19] Daw M S and Baskes M I 1984 Embedded-atom method: Derivation and application to impurities, surfaces, and other defects in metals *Phys. Rev. B* **29** 6443–53

[20] Daw M S 1989 Model of metallic cohesion: The embedded-atom method *Phys. Rev. B* **39** 7441–52

[21] Daw M S, Foiles S M and Baskes M I 1993 The embedded-atom method: a review of theory and applications *Mater. Sci. Rep.* **9** 251–310

[22] Elliott R S and Tadmor E B 2016 Kim application programming interface (api) standard `https://openkim.org/kim-api/`

[23] Asap – as soon as possible molecular dynamics calculator for ase `https://wiki.fysik.dtu.dk/asap/`

[24] Smith W and Forester T R 1996 DL_POLY_2.0: A general-purpose molecular dynamics simulation package *J. Mol. Graph.* **14** 136–41

[25] 2015 Dl_poly molecular simulation package `http://www.stfc.ac.uk/SCD/44516.aspx`

[26] Gale J D 1997 GULP: A computer program for the symmetry-adapted simulation of solids *J. Chem. Soc.-Faraday Trans.* **93** 629–37

[27] 2015 Gulp website `http://nanochemistry.curtin.edu.au/gulp/`

[28] Stadler J, Mikulla R and Trebin H R 1997 IMD: A software package for molecular dynamics studies on parallel computers *Int. J. Mod. Phys. C* **8** 1131–40

[29] Roth J 2013 *Sustained Simulation Performance 2013* chapter IMD: A Typical Massively Parallel Molecular Dynamics Code for Classical Simulations – Structure, Applications, Latest Developments Springer International Publishing, Cham 63–76

[30] 2015 Imd: The itap molecular dynamics program `http://imd.itap.physik.uni-stuttgart.de`

[31] Plimpton S 1995 Fast parallel algorithms for short-range molecular dynamics *J. Comp. Phys.* **117** 1–19

[32] 2015 Large-scale atomic/molecular massively parallel simulator (lammps) `http://lammps.sandia.gov`

[33] libatoms+quip website `http://www.libatoms.org`

[34] Tadmor E B, Ortiz M and Phillips R 1996 Quasicontinuum analysis of defects in solids *Philos. Mag. A* **73** 1529–63

[35] Nissen-Hooper J, Karls D S, Elliott R S and Tadmor E B 2016 Visualization of material properties in the openkim project In preparation

[36] Bazant M Z and Kaxiras E 1996 Modeling of covalent bonding in solids by inversion of cohesive energy curves *Phys. Rev. Lett.* **77** 4370

[37] Bazant M Z, Kaxiras E and Justo J 1997 Environment-dependent interatomic potential for bulk silicon *Phys. Rev. B* **56** 8542

[38] Justo J F, Bazant M Z, Kaxiras E, Bulatov V and Yip S 1998 Interatomic potential for silicon defects and disordered phases *Phys. Rev. B* **58** 2539

[39] Wen M, Whalen S M, Elliott R S and Tadmor E B 2015 Interpolation effects in tabulated interatomic potentials *Modell. Simul. Mater. Sci. Eng.* **23** 074008

[40] Mishin Y, Mehl M and Papaconstantopoulos D 2005 Phase stability in the fe–ni system: Investigation by first-principles calculations and atomistic simulations *Acta Mater.* **53** 4029–41

[41] Tersoff J 1986 New empirical model for the structural properties of silicon *Phys. Rev. Lett.* **56** 632

[42] Tersoff J 1989 Modeling solid-state chemistry: Interatomic potentials for multicomponent systems *Phys. Rev. B* **39** 5566

[43] Tangney P and Scandolo S 2002 An ab initio parametrized interatomic force field for silica *J. Chem. Phys.* **117**

8898–904
[44] Powell M J D 1965 A method for minimizing a sum of squares of non-linear functions without calculating derivatives *Comp. J.* **7** 303–7
[45] Karls D S 2012 Original EDIP potential for silicon https://openkim.org/cite/MO_958932894036_001
[46] Karls D S 2012 A C-based implementation of the EDIP three-body bond-order potential of Bazant and Kaxiras https://openkim.org/cite/MD_506186535567_001
[47] Bierbaum M, Alemi A, Karls D S, Wennblom T J, Elliott R S, Sethna J P and Tadmor E B 2016 The OpenKIM Processing Pipeline: Avirtual machine cloud-based automatic materials property computation engine In preparation
[48] Karls D S 2016 *Transferability of Empirical Potentials and the Knowledgebase of Interatomic Models (KIM)* Ph.D. thesis University of Minnesota
[49] Transtrum M K, Machta B B and Sethna J P 2010 Why are nonlinear fits to data so challenging? *Phys. Rev. Lett.* **104** 060201
[50] Transtrum M K and Sethna J P 2012 Improvements to the Levenberg–Marquardt algorithm for nonlinear least-squares minimization *arXiv preprint arXiv:1201.5885*
[51] Transtrum M K and Sethna J P 2012 Geodesic acceleration and the small-curvature approximation for nonlinear least squares *arXiv preprint arXiv:1207.4999*
[52] Madsen K, Nielsen H B and Tingleff O 2004 Methods for non-linear least squares problems
[53] Levenberg K 1963 A method for the solution of certain problems in least squares *Quart. Appl. Math.* **2** 168–4
[54] Marquardt D W 1963 An algorithm for least-squares estimation of nonlinear parameters *J. Soc. Ind. Appl. Math.* **11** 431–41
[55] Transtrum M K 2015 Geodesic Levenberg-Marquart code package version 1.1 https://sourceforge.net/projects/geodesiclm/
[56] Broyden C G 1965 A class of methods for solving nonlinear simultaneous equations *Mathematics of computation* **19** 577–93
[57] Frederiksen S L, Jacobsen K W, Brown K S and Sethna J P 2004 Bayesian ensemble approach to error estimation of interatomic potentials *Phys. Rev. Lett.* **93** 165501
[58] Lourakis M Jul. 2004 levmar: Levenberg–Marquardt nonlinear least squares algorithms in C/C++ http://www.ics.forth.gr/~lourakis/levmar/
[59] Novoselov K S, Jiang D, Schedin F, Booth T J, Khotkevich V V, Morozov S V and Geim A K 2005 Two-dimensional atomic crystals *Proc. Natl. Acad. Sci. USA* **102** 10451–3

arXiv:1611.03916v1 [cond-mat.other] 12 Nov 2016

Supplementary information for "A KIM-compliant *potfit* for fitting sloppy interatomic potentials: Application to the EDIP model for silicon"

Mingjian Wen, Junhao Li, Peter Brommer, Ryan S. Elliott, James P. Sethna and Ellad B. Tadmor‡

E-mail: tadmor@umn.edu
Department of Aerospace Engineering and Mechanics, University of Minnesota, Minneapolis, MN 55455, USA
[1] Warwick Centre for Predictive Modelling, School of Engineering, and Centre for Scientific Computing, University of Warwick, Coventry CV4 7AL, UK

‡ Corresponding author.

1. The EDIP potential for silicon

In this section we describe the functional form of the Environment Dependent Interatomic Potential (EDIP) for bulk silicon developed by Bazant et al. [1–3]. The total potential energy of a configuration of n atoms is assumed to take the form $E = \sum_{i=1}^{n} E_i$, where the contribution to the energy of atom i is given by

$$E_i = \sum_{j \neq i} V_2(R_{ij}, Z_i) + \sum_{j \neq i} \sum_{k \neq i, k > j} V_3(\vec{R}_{ij}, \vec{R}_{ik}, Z_i). \tag{1}$$

The "local environmental dependence" of atom i is introduced through its effective coordination number, defined by,

$$Z_i = \sum_{m \neq i} f(R_{im}), \tag{2}$$

where $f(R_{im})$ is a cutoff function given by,

$$f(r) = \begin{cases} 1 & \text{if } r < c \\ \exp\left(\dfrac{\alpha}{1 - x^{-3}}\right) & \text{if } c < r < a \\ 0 & \text{if } r > a \end{cases} \tag{3}$$

in which $x = (r - c)/(a - c)$.

The two-body term V_2 includes a repulsive part and an attractive part:

$$V_2(r, Z) = A \left[\left(\frac{B}{r} \right)^{\rho} - p(Z) \right] \exp\left(\frac{\sigma}{r - a} \right) \tag{4}$$

with the attractive bond-order function taking the form $p(Z) = e^{-\beta Z^2}$.

The three-body term V_3 takes the form of a product of radial and angular functions:

$$V_3(\vec{R}_{ij}, \vec{R}_{ik}, Z_i) = g(R_{ij}) g(R_{ik}) h(l_{ijk}, Z_i). \tag{5}$$

The radial contribution to V_3 is given by $g(r) = \exp(\gamma/(r - a))$, and the angular contribution is

$$h(l, Z) = \lambda \left[\left(1 - e^{-Q(Z)(l + \tau(Z))^2} \right) + \eta Q(Z) \left(l + \tau(Z) \right)^2 \right] \tag{6}$$

with the bond angle dependence contained in $l_{ijk} = \cos\theta_{ijk} = \vec{R}_{ij} \cdot \vec{R}_{ik} / R_{ij} R_{ik}$ and the coordination number dependence controlled by the function $Q(Z) = Q_0 e^{-\mu Z}$. The function $\tau(Z)$, which controls the equilibrium angle of the three-body interaction, is chosen to have the form

$$\tau(Z) = u_1 + u_2 \left(u_3 e^{-u_4 Z} - e^{-2u_4 Z} \right), \tag{7}$$

where $u_1 = -0.165799$, $u_2 = 32.557$, $u_3 = 0.286198$ and $u_4 = 0.66$ are parameters chosen such that $\tau(Z)$ yields correct theoretical bond angle at $Z = 2, 3, 4$, and 6, and interpolates smoothly between them.

The EDIP potential has 13 adjustable parameters: A, B, ρ, β, σ, λ, η, γ, μ, α, Q_0, a, and c, and their values are listed in table 1. The thresholds values a and c used in the cutoff function $f(r)$ are fixed during the fitting process. Consequently, only 11 parameters are optimized.

Table 1. EDIP Parameter values for silicon

Parameter	Value
A	7.9821730 eV
B	1.5075463 Å
ρ	1.2085196
β	0.0070975
σ	0.5774108 Å
λ	1.4533108 eV
η	0.2523244
γ	1.1247945 Å
μ	0.6966326
α	3.1083847
Q_0	312.1341346
a	3.1213820 Å
c	2.5609104 Å

2. Training set

The training set consists of the energy and forces computed with EDIP for a periodic configuration of 1000 atoms perturbed off the ideal diamond structure positions. The training set file "supplementary_traning_set.txt" accompanies this article. The header of the file specifies the supercell lattice vectors and the average energy per atom. This is followed by a list giving for each atom the atomic species code, atom coordinates, and force on the atom. See the *potfit* webpage http://www.potfit.net/wiki/doku.php for details on the format of the training set file.

References

[1] Martin Z Bazant and Efthimios Kaxiras. Modeling of covalent bonding in solids by inversion of cohesive energy curves. *Phys. Rev. Lett.*, 77(21):4370, 1996.

[2] Martin Z Bazant, Efthimios Kaxiras, and Joao F Justo. Environment-dependent interatomic potential for bulk silicon. *Phys. Rev. B*, 56(14):8542, 1997.

[3] Joao F Justo, Martin Z Bazant, Efthimios Kaxiras, Vasily V Bulatov, and Sidney Yip. Interatomic potential for silicon defects and disordered phases. *Phys. Rev. B*, 58(5):2539, 1998.